



PROJECT AIR FORCE

CHILD POLICY

CIVIL JUSTICE

EDUCATION

ENERGY AND ENVIRONMENT

HEALTH AND HEALTH CARE

INTERNATIONAL AFFAIRS

NATIONAL SECURITY

POPULATION AND AGING

PUBLIC SAFETY

SCIENCE AND TECHNOLOGY

SUBSTANCE ABUSE

TERRORISM AND
HOMELAND SECURITY

TRANSPORTATION AND
INFRASTRUCTURE

This PDF document was made available from www.rand.org as a public service of the RAND Corporation.

[Jump down to document ▼](#)

The RAND Corporation is a nonprofit research organization providing objective analysis and effective solutions that address the challenges facing the public and private sectors around the world.

Support RAND

[Purchase this document](#)

[Browse Books & Publications](#)

[Make a charitable contribution](#)

For More Information

Visit RAND at www.rand.org

Explore [RAND Project AIR FORCE](#)

View [document details](#)

Limited Electronic Distribution Rights

This document and trademark(s) contained herein are protected by law as indicated in a notice appearing later in this work. This electronic representation of RAND intellectual property is provided for non-commercial use only. Permission is required from RAND to reproduce, or reuse in another form, any of our research documents.

| Report Documentation Page | | | | Form Approved OMB No. 0704-0188 | |
|--|------------------------------------|-------------------------------------|----------------------------|--|---------------------------------|
| Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. | | | | | |
| 1. REPORT DATE 2005 | | 2. REPORT TYPE | | 3. DATES COVERED - | |
| 4. TITLE AND SUBTITLE Aging Aircraft Repair-Replacement Decisions with Depot-Level Capacity as a Policy Choice Variable | | | | 5a. CONTRACT NUMBER | |
| | | | | 5b. GRANT NUMBER | |
| | | | | 5c. PROGRAM ELEMENT NUMBER | |
| 6. AUTHOR(S) | | | | 5d. PROJECT NUMBER | |
| | | | | 5e. TASK NUMBER | |
| | | | | 5f. WORK UNIT NUMBER | |
| 7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Department of the Air Force, Strategic Planning Division, Directorate of Plans, Washington, DC, 20330 | | | | 8. PERFORMING ORGANIZATION REPORT NUMBER | |
| 9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) | | | | 10. SPONSOR/MONITOR'S ACRONYM(S) | |
| | | | | 11. SPONSOR/MONITOR'S REPORT NUMBER(S) | |
| 12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited | | | | | |
| 13. SUPPLEMENTARY NOTES The original document contains color images. | | | | | |
| 14. ABSTRACT see report | | | | | |
| 15. SUBJECT TERMS | | | | | |
| 16. SECURITY CLASSIFICATION OF: | | | 17. LIMITATION OF ABSTRACT | 18. NUMBER OF PAGES 71 | 19a. NAME OF RESPONSIBLE PERSON |
| a. REPORT unclassified | b. ABSTRACT unclassified | c. THIS PAGE unclassified | | | |

This product is part of the RAND Corporation monograph series. RAND monographs present major research findings that address the challenges facing the public and private sectors. All RAND monographs undergo rigorous peer review to ensure high standards for research quality and objectivity.

Aging Aircraft Repair-Replacement Decisions with Depot-Level Capacity as a Policy Choice Variable

Edward G. Keating, Don Snyder,
Matthew Dixon, Elvira N. Loredó

Prepared for the United States Air Force
Approved for public release; distribution unlimited



PROJECT AIR FORCE

The research reported here was sponsored by the United States Air Force under contract F49642-01-C-0003. Further information may be obtained from the Strategic Planning Division, Directorate of Plans, Hq USAF.

Library of Congress Cataloging-in-Publication Data

Aging aircraft repair-replacement decisions with depot-level capacity as a policy choice
variable / Edward G. Keating ... [et al.].

p. cm.

Includes bibliographical references.

"MG-241."

ISBN 0-8330-3692-0 (pbk. : alk. paper)

1. Airplanes, Military—United States. 2. Armed forces—Appropriations and expenditures. 3. Product life cycle. I. Keating, Edward G., 1965– II. Title.

TL685.3.A4 2005

358.4'183'0973—dc22

2004023486

The RAND Corporation is a nonprofit research organization providing objective analysis and effective solutions that address the challenges facing the public and private sectors around the world. RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

RAND® is a registered trademark.

© Copyright 2005 RAND Corporation

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from RAND.

Published 2005 by the RAND Corporation

1776 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138

1200 South Hayes Street, Arlington, VA 22202-5050

201 North Craig Street, Suite 202, Pittsburgh, PA 15213-1516

RAND URL: <http://www.rand.org/>

To order RAND documents or to obtain additional information, contact

Distribution Services: Telephone: (310) 451-7002;

Fax: (310) 451-6915; Email: order@rand.org

Preface

The decision on whether to maintain or replace an aging system is a common one. Anyone who owns an automobile, for instance, eventually grapples with this issue. At some point, it seems wrong to “throw good money after bad” and continue to repair an aging system. But replacement systems typically entail considerable up-front investment.

This monograph continues a sequence of RAND Project AIR FORCE reports bearing on aging aircraft and the replacement-repair decision, of which Greenfield and Persselin (2002), Pyles (2003), and Keating and Dixon (2003) are recent examples. In this study, we present a model of the repair-replacement decision and data describing the C-5A cargo aircraft. We show how our methodology can be used to assess the desirability of a proposed large C-5A aircraft modification. We also present a new methodology to explore the benefits of increases in depot-level capacity.

This study was conducted for a project entitled “Understanding and Addressing the Effects of Aging Aircraft,” sponsored by AF/ILM and AF/XPX and conducted within the Resource Management Program of RAND Project AIR FORCE.

This research is intended to be of interest to U.S. Air Force and other Department of Defense acquisition and logistics personnel.

RAND Project AIR FORCE

RAND Project AIR FORCE (PAF), a division of the RAND Corporation, is the U.S. Air Force's federally funded research and

development center for studies and analyses. PAF provides the Air Force with independent analyses of policy alternatives affecting the development, employment, combat readiness, and support of current and future aerospace forces. Research is conducted in four programs: Aerospace Force Development; Manpower, Personnel, and Training; Resource Management; and Strategy and Doctrine.

Additional information about PAF is available on our web site at <http://www.rand.org/paf>.

Contents

Preface iii

Figures vii

Tables ix

Summary xi

Acknowledgments xv

Abbreviations xvii

Glossary xix

CHAPTER ONE

Introduction 1

CHAPTER TWO

Modeling the Decision to Repair or Replace an Aging Aircraft 3

CHAPTER THREE

The C-5A Modification Versus Replacement Decision 7

Background 7

C-5A and C-17 Costs 11

Cost Factor Growth Rates 13

Aircraft Availability 17

Optimal C-5A Retirement Calculation 18

Robustness Explorations 19

Discussion 24

CHAPTER FOUR

Consideration of C-5A Depot-Level Capacity 27

Growing Hands-On Time During PDM and Repair 36

Investing in More Depot-Level Capacity 39

CHAPTER FIVE

Conclusions 43

References 45

Figures

| | | |
|-------|---|----|
| 3.1. | The C-5 Aircraft | 8 |
| 3.2. | The C-17 Aircraft | 10 |
| 3.3. | C-5A On- and Off-Equipment Maintenance Hours Per Possessed Hour | 14 |
| 3.4. | C-5A PDM Hours Per Aircraft | 15 |
| 3.5. | C-5A MC and Composite Availability Rates | 17 |
| 3.6. | C-5A Optimal Replacement Baseline Case | 19 |
| 3.7. | C-5A Optimal Replacement as a Function of When the Median C-5A Gets the RERP | 21 |
| 3.8. | C-5A Optimal Replacement as a Function of the Number of C-17s Required | 22 |
| 3.9. | C-5A Optimal Replacement as a Function of the Composite Availability Gain After the Age 44 Mod | 23 |
| 3.10. | C-5A Optimal Replacement as a Function of the Aircraft Overhaul Cost Growth Rate Past Age 31 | 24 |
| 4.1. | Lower Bound on C-5A Operating Command Possession Rate as a Function of Depot-Level Capacity | 36 |
| 4.2. | Lower Bound on C-5A Operating Command Possession Rate as a Function of R_0 | 37 |
| 4.3. | Required Level of C-5A Capacity as a Function of Minimum Hands-On Time, R_0 | 38 |
| 4.4. | Annualized Cost Per Available Year with Investment in Extra C-5A Depot-Level Capacity | 40 |
| 4.5. | C-5A Possession Rate With and Without Extra Depot-Level Capacity | 41 |

Tables

3.1. Comparing the C-5 and the C-17 10

3.2. AFTOC Annual Per-Active-Aircraft Cost Factors (FY04 \$) . . . 11

3.3. Annual Sustainment Costs and Assumed Growth Rates of
the C-5A and C-17 Fleets (per aircraft) 16

3.4. Estimated C-5A Costs 20

Summary

In Keating and Dixon (2003), we presented a model for determining when it would be optimal to retire, rather than continue to repair, an aging system. This work extends Keating and Dixon along two dimensions.

First, we extend our methodology to examine whether a proposed modification (mod) is worthwhile relative to retiring an aircraft.

Second, we develop a new methodology to explore the desirability of additional investment in depot-level capacity.

Modeling the Decision to Repair or Replace an Aging Aircraft

In our model, we consider the discounted expenditure and availability flows emanating from repairing versus replacing an aging system.

If one assumes the Air Force's goal is to minimize its average cost per available year (or day), the Air Force should keep an incumbent aircraft until its incremental cost per available year exceeds the average cost per available year of a replacement system. (See pp. 5–6.)

This approach should be used prospectively. For example, one estimates ahead of time when it is thought the optimum will be achieved and has a replacement system prepared to enter service at that time.

The C-5A Modification Versus Replacement Decision

We show how our repair-replace methodology can be used to assess the desirability of modifying versus retiring an aging aircraft. We use data on the C-5A to illustrate our methodology.

The C-5 is the Air Force's largest cargo-carrying aircraft. There are currently 72 C-5As; the fleet's average age is somewhat over 30 years. The Air Force plans to retire 12 C-5As by the end of fiscal year (FY) 2006, leaving a fleet of 60 C-5As.

The C-5A fleet is scheduled to receive a major mod or Reliability Enhancement and Re-engining Program (RERP) sometime in the 2010 decade, requiring work on each aircraft's airframe, avionics, engines, landing gear, and other equipment. Tirpak (2004) cites the RERP's cost as \$75 million per aircraft, so the Air Force needs to carefully consider the C-5A's future before making this commitment. If, instead, the C-5A fleet is retired in front of the mod, the Air Force could purchase additional C-17s or develop an alternative aircraft. In this study, we focus on the C-17 as a replacement for the C-5A because of the availability of data on the costs of C-17s.

To demonstrate our model, we used Air Force Total Ownership Cost (AFTOC) factors for the C-5 and the C-17. We note with concern, however, that AFTOC factors do not appropriately differentiate between C-5As and C-5Bs for our analytic purposes.

Pyles (2003) presents Air Force data that show considerable growth in the C-5A's programmed depot maintenance (PDM) package. On the other hand, the C-5A has not seen adverse trends in its on- and off-equipment maintenance costs. The C-5A has experienced only a minor decline in the fraction of the fleet that is possessed by operating commands and is mission capable (which we label the composite availability rate).

With our baseline assumptions, our model indicates that it is not optimal to undertake the proposed C-5A mod in 2015. However, the mod appears to be worthwhile if it can occur earlier, perhaps in 2010. As the mod is delayed, C-5A performance degrades and costs increase to the point, eventually, that it is not worthwhile to undertake the mod. (See pp. 8–21.)

We then undertook a robustness exploration to see which other parameters were critical in undergirding our findings. Along with date of the RERP, another key parameter is the number of C-17s it would take to acceptably replace 60 C-5As. As the number of required C-17s increases, the desirability of the C-5A RERP increases. If 70 or more C-17s are required, the 2015 C-5A RERP is worthwhile, given our other parameters. (See pp. 21–22.)

The Air Force is currently implementing the RERP on two C-5Bs and a C-5A. The virtue of undertaking the RERP on a few aircraft is that it will shed light on the magnitude of the availability gains from the project. On the other hand, our model indicates that delay in the C-5A RERP tends to diminish its desirability.

We view our C-5A findings as illustrative and suggestive, rather than definitive, particularly in light of concerns with our C-5A cost parameters.

Consideration of C-5A Depot-Level Capacity

We extend our methodology to assess the desirability of additional investment in depot-level capacity (including additional facilities, repair equipment, labor, and spare parts). We do not assess the optimal form of such investment.

We know the Air Force values having aircraft available to its operating commands. Our repair-replace calculation assumes the Air Force would be willing to someday pay the likely high costs of a replacement system. Given a belief (or inferred preference) about how much an available aircraft year is worth, we can calibrate the desirability of increases in depot-level capacity that place aircraft in operators' hands faster.

Queuing is a source of considerable delay in depot-level maintenance. Hence, we developed a model of the programmed depot-maintenance process as a multinode or multistep closed network.

Using our depot-level queuing model, we consider the effect of increasing capacity (e.g., purchasing more repair equipment, spare parts, or hiring more workers). When capacity increases, the rate at

which operating commands possess aircraft increases as aircraft spend less time in depot-level queues.

Intriguingly, our model, using illustrative C-5A data, suggests that, while an increase in depot-level capacity is desirable, it does not delay the optimal retirement of the aircraft. Instead, the effect is to increase the operating command's possession rate while the aircraft is used. (See pp. 40–42.)

We think our findings are explained by the inferred preference exercise upon which this analysis is built: We assume the Air Force will eventually replace the C-5A with the C-17, a nontrivially expensive aircraft. Hence, we infer that the Air Force places high valuation on an available C-17, and hence C-5A, year. Given this inferred preference, the model is then averse to having large-scale queuing and delay in the C-5A depot maintenance system. Given the inferred value of these aircraft, it is not reasonable to have them wait in many or long depot-level queues.

Acknowledgments

Lieutenant Colonel Randall Burke, AF/ILMY, was our helpful project monitor. Crash Lively of Battelle assisted us with interpreting Air Force Total Ownership Cost (AFTOC) data.

We received constructive reviews of this report from Professor Michael Alles of Rutgers University and from our colleague Isaac Porche. This research was presented at the Western Economic Association annual conference in Vancouver, British Columbia, on July 1, 2004. We appreciate the thought-provoking insights of our discussant, David Trybula, of the Office of the Secretary of Defense.

We also appreciate the assistance of our colleagues Laura Baldwin, Jean Gebman, Gregory Hildebrandt, Michael Kennedy, Ray Pyles, and C. Robert Roll on this study. Jane Siegel and Alissa Hiraga helped with the production of this document and Jeanne Heller edited it.

Of course, the authors alone are responsible for any errors that may remain in the document.

Abbreviations

| | |
|-------|--|
| AFTOC | Air Force Total Ownership Cost |
| FY | Fiscal Year |
| MC | Mission Capable |
| OMB | Office of Management and Budget |
| PDM | Programmed Depot Maintenance |
| REMIS | Reliability and Maintainability Information System |
| RERP | Reliability Enhancement and Re-engining Program |

Glossary

Composite availability—The product of an aircraft's mission-capable rate and the fraction of the fleet currently held by operating commands, i.e., not in the depot system. The composite availability rate is the fraction of the total fleet held by operating commands and mission capable. We label time t 's composite availability rate $Availability_t$. Composite availability can be computed at a point in time or over a period of time, such as a month or a year.

Contractor logistical support—A contractor (Boeing in the case of the C-17) provides much of the depot-level (and possibly other types of) maintenance on an aircraft as opposed to having government-owned depots (e.g., Tinker Air Force Base, Robins Air Force Base, Hill Air Force Base) provide such maintenance.

Cost per available year—The appropriately discounted life-cycle total cost of a system divided by the appropriately discounted sum of expected fractions of years an aircraft will be mission capable and possessed by operating commands. Cost per available year can also be termed annualized cost.

Example: Suppose an aircraft has expected total, inflation-adjusted costs ($Expend_t$) of \$10 million this year, \$11 million next year, and \$8 million two years from now. Suppose, too, that we expect an 80 percent composite availability rate this year, 75 percent next year, and 50 percent two years from now. With a 3.5 percent real discount rate, the system's cost per available year would be

$$\frac{10,000,000 + \frac{11,000,000}{1.035} + \frac{8,000,000}{1.035^2}}{0.8 + \frac{0.75}{1.035} + \frac{0.5}{1.035^2}} = 14,108,769.$$

Cost per flying hour—Total yearly costs divided by total yearly flying hours.

Depot-level capacity—The capacity to perform required repairs on aircraft in the depot system. More capacity implies aircraft should spend less time (or get more done while) in the depot system. We use this term as a general rubric. It includes, but is not limited to, facilities, equipment, labor, and spare parts used in performing depot-level activities.

Discount rate—The rate at which future costs and benefits are discounted relative to today. A positive discount rate reflects a preference for having cash or getting a benefit today relative to next year. For calendar year 2004, the Office of Management and Budget prescribes use of a 3.5 percent real (already adjusted for expected inflation) discount rate.

Mission-capable rate—The percentage of a fleet thought to be able to safely perform its missions. An aircraft can be mission capable without being fully mission capable. For instance, an aircraft may be capable of performing a daytime mission even with broken night-vision equipment.

Mod—A large-scale modification of an existing aircraft, usually to enhance capability, improve maintainability, or address safety concerns. The proposed mod on the C-5 is called the Reliability Enhancement and Re-engining Program (RERP).

Off-equipment maintenance—Maintenance performed at an operating base's backshop. A broken part is removed from an aircraft and brought to an on-base facility where it is repaired and then returned either to the aircraft or to inventory.

On-equipment maintenance—Maintenance performed directly on the aircraft at an operating base. On-equipment and off-

equipment maintenance combine to form organizational maintenance.

Possession rate—The fraction of the time an aircraft is held by an operating command (e.g., Air Mobility Command) rather than being in the depot system.

Programmed depot maintenance (PDM)—A major aircraft overhaul performed at either a government-owned (organic) or private-sector depot facility. Programmed depot maintenance occurs on a schedule, such as five years after the last PDM visit for the C-5A. PDM can also be termed aircraft overhaul or heavy maintenance. In addition, aircraft may need unscheduled depot maintenance (e.g., a bird strike damages an aircraft and sophisticated depot-level maintenance is required).

Queuing—When aircraft have to wait while in the depot system. Queues may result from not having enough facilities, equipment, skilled labor, and/or required spare parts. Of course, it is probably not cost effective to have so much depot-level capacity that no queues ever occur.

R_0 —The minimum hands-on flow time required to perform programmed depot maintenance. It is equivalent to the time required to complete PDM if the aircraft did not have to wait in queues.

Introduction

In earlier work (Keating and Dixon, 2003), we introduced an approach for determining when it is optimal to replace an aging system. Under our procedure, one keeps an existing system until that system's incremental costs exceed the average cost of a replacement system. In Keating and Dixon (2003), we applied our computational approach to the C-21A transport aircraft and the KC-135 tanker aircraft. Based on available data and our assumptions, the model indicated that it would be optimal to operate the C-21A until sometime around 2020, while it would be optimal to replace the KC-135 fleet sometime late in this decade.

In this report, we extend our analysis approach along two dimensions. We discuss how our modeling approach can be used to evaluate potential modifications (mods) versus the alternative of retiring an aircraft. As we will show, our modeling approach accommodates large expenditure spikes and provides insight as to whether a mod should be undertaken or whether an aircraft should be retired before the mod occurs.

We then apply the model to the C-5A cargo aircraft, because the fleet is scheduled to receive an expensive mod in the next decade or so. We assess the desirability of the proposed C-5A mod and then extend our modeling tool in a different direction. Specifically, we experiment with making depot-level capacity a choice variable. Additional capacity could take the form of new physical equipment or facilities, but it also could be extra labor or spare parts. The virtue of

additional depot-level capacity is that it allows aircraft to spend less time in the depot system and hence more time possessed by the operating commands.

We rely on an inferred preference calculation. We assert that the Air Force is willing to someday invest in a replacement system, so that replacement system's cost per available year can form a lower bound for how much the Air Force values the availability of an existing system.

Given our parameter estimates and the assumptions underlying our new methodology, we estimate that there could be considerable value in increased C-5A depot-level capacity. It is not desirable to have a highly valued system waiting in many or lengthy depot-level queues. However, we have not investigated the best way to obtain extra capacity, such as more or better labor, additional spare parts, more repair equipment or facilities.

In Chapter Two, we sketch our basic model, reprising material from Keating and Dixon (2003). In Chapter Three, we apply the model to the C-5A, examining the desirability of that aircraft's planned upcoming mod. In Chapter Four, we again use C-5A data, but in a more exploratory vein, to illustrate how increased depot-level capacity might be evaluated. Chapter Five gives our conclusions.

Modeling the Decision to Repair or Replace an Aging Aircraft

Consider a type of aircraft (e.g., a cargo aircraft) that the Air Force envisions having in its inventory (in one form or another) into the foreseeable future. This string of future aircraft would have a string of discounted future expenses

$$x = \sum_{t=1}^{\infty} \frac{\textit{Expend}_{Rt}}{(1 + \textit{Discount})^{t-1}}$$

where the first replacement aircraft (we use R to denote the replacement, rather than incumbent, aircraft) flew in Year 1 and $\textit{Discount}$ is the Office of Management and Budget's prescribed real interest rate, e.g., 3.5 percent in calendar year 2004.

These future aircraft would provide a string of available days (or fractions of years). The Office of Management and Budget's Circular A-94 prescribes that "all future benefits and costs, including non-monetized benefits and costs, should be discounted." Hence, symmetric to x , we can define a future availability sum

$$y = \sum_{t=1}^{\infty} \frac{\textit{Availability}_{Rt}}{(1 + \textit{Discount})^{t-1}}.$$

Consider, then, an existing incumbent aircraft, I . Without loss of generality, we will focus on the question of whether the Air Force should fly the incumbent aircraft one more year and then replace it or replace it right now. (There is no loss of generality in this simplification. If, for instance, it is actually optimal to operate the incumbent aircraft for six more years, it will certainly be optimal to operate it for one more year.)

If the incumbent aircraft is kept only one more year, the Air Force spends $Expend_{I1}$ to sustain and operate the aircraft and receives fraction $Availability_{I1}$ years of availability (or $365 * Availability_{I1}$ days). $Expend_{I1}$ would include, for instance, the maintenance, fuel, and labor costs associated with flying the incumbent aircraft another year. Then, starting next year, the replacement's strings of expenses and availability start. The replacement's expenses would include production, testing, and research and development costs not yet borne, as well as the operating costs of the new aircraft. Therefore, the discounted infinite sum of expenditures associated with keeping the incumbent one more year would be

$$Expend_{I1} + \frac{x}{1 + Discount},$$

while the sum of availability would be

$$Availability_{I1} + \frac{y}{1 + Discount}.$$

(The Air Force still pays x to get y years of availability from the replacements, but it has been pushed back by one year and hence is discounted.)

Suppose the Air Force's sole goal is to minimize its expenditures. Then it would operate the existing aircraft another year if and only if

$$Expend_{I1} + \frac{x}{1 + Discount} \leq x \text{ or } Expend_{I1} \leq \frac{x * Discount}{1 + Discount}.$$

However, we do not think the Air Force's sole goal is expenditure minimization because this goal would not consider the availability of its aircraft. Instead, the Air Force's goal may be to minimize its average cost per available year (or day). With this metric, retaining the incumbent aircraft for one more year results in an average cost per available year of

$$\frac{\text{Expend}_{I1} + \frac{x}{1 + \text{Discount}}}{\text{Availability}_{I1} + \frac{y}{1 + \text{Discount}}}.$$

Therefore, the Air Force should repair, rather than replace, an aging system for one more year if and only if

$$\frac{\text{Expend}_{I1} + \frac{x}{1 + \text{Discount}}}{\text{Availability}_{I1} + \frac{y}{1 + \text{Discount}}} \leq \frac{x}{y} \text{ or } \frac{\text{Expend}_{I1}}{\text{Availability}_{I1}} \leq \frac{x}{y}.$$

If one wished to apply this optimality condition, we would not recommend waiting until the optimal retirement year to take action because there are lags associated with acquiring a replacement system. Instead, this approach should be used prospectively (e.g., estimating ahead of time when it is thought the optimum will be achieved and having a replacement system prepared to enter service at that time).

This modeling approach can accommodate a number of real-world phenomena. For example, the desirability of a mod that involves a one-time expenditure spike can be assessed by comparing costs per available year immediately preceding the mod with those associated with the optimal retirement year conditional on the mod occurring.

Let B equal the number of years before the mod is required and M^* be the optimal retirement date conditional on the mod occur-

ring. Then the cost per available year if the mod does not occur and the aircraft is retired in B years is

$$AC_{NM} = \frac{\sum_{t=1}^B \frac{Expend_{I_t}}{(1 + Discount)^{t-1}} + \frac{x}{(1 + Discount)^B}}{\sum_{t=1}^B \frac{Availability_{I_t}}{(1 + Discount)^{t-1}} + \frac{y}{(1 + Discount)^B}},$$

while the cost per available year if the mod does occur is

$$AC_M = \frac{\sum_{t=1}^{M^*} \frac{Expend_{I_t}}{(1 + Discount)^{t-1}} + \frac{x}{(1 + Discount)^{M^*}}}{\sum_{t=1}^{M^*} \frac{Availability_{I_t}}{(1 + Discount)^{t-1}} + \frac{y}{(1 + Discount)^{M^*}}}.$$

Of course, AC_M incorporates both the up-front costs and availability benefits associated with the mod program. The mod should be undertaken if and only if $AC_M < AC_{NM}$. We present an example of a mod desirability calculation in Chapter Three.

In Chapter Four, we further enhance the model by making the amount of depot-level capacity a choice variable.

In Keating and Dixon (2003), we applied this approach to the C-21A transport and KC-135 tanker aircraft. Here, we illustrate the model and its enhancements using data on the C-5A cargo aircraft. As in our earlier work, we use C-5A data to describe how our methodology works rather than to definitively assert what the Air Force should do with this system.

The C-5A Modification Versus Replacement Decision

In this chapter, we show how our repair-replace methodology can be used to assess the desirability of modifying versus retiring an aging aircraft. We use data on the C-5A to illustrate our methodology.

Background

The C-5A and its descendant, the C-5B, are large cargo-transport aircraft. Figure 3.1 shows a C-5.

Lockheed-Georgia was awarded the original C-5A contract in 1965. Not long into production, stress problems caused significant cracks in the wings. At the same time, problems with the main landing gear were discovered in testing. There was considerable acquisition cost escalation. See Pike (2001).

According to the Air Force's Reliability and Maintainability Information System (REMIS), of the 72 C-5As currently operating, the Air Force accepted its oldest C-5A on May 26, 1969, and its newest C-5A on November 4, 1975. The Air Force plans to retire 12 C-5As by the end of fiscal year (FY) 2006. There are 50 newer C-5Bs.

In January 1998, while undergoing periodic depot maintenance at Kelly Air Force Base, Texas, cracks were noticed in a C-5A's tail section. Further inspection revealed other C-5As with similar cracks.

Figure 3.1
The C-5 Aircraft



RAND MG241-3.1

Photo from <http://www.af.millfactsheets/factsheet.asp?fsID=84>.

In response to this problem, a variety of modifications have been undertaken to C-5A tails. See Pike (2001).

The C-5A is facing a planned major modification (termed the Reliability Enhancement and Re-engining Program, or RERP). The RERP will give each aircraft four new engines and pylons, as well as a new engine-driven generator and auxiliary power unit. Additional improvements are to be made to the landing gear, hydraulics, flight controls, and environmental control system. The RERP will cost an estimated \$75 million per aircraft. See Tirpak (2004). The RERP is scheduled to begin with the C-5Bs and then proceed to the C-5As in the 2012–2018 timeframe. We assume the median C-5A will obtain its RERP in 2015 when it is 44 years old.

A key policy question for the Air Force is whether this investment is worth undertaking for the remaining 60 C-5As or whether, alternatively, the aircraft should be retired in front of the RERP. (We have not analyzed the possibility of keeping the C-5A fleet without performing this mod.) In July 2004, the Air Force's Fleet Viability

Board concluded that the C-5A fleet has at least 25 years of service life remaining. See Jablonski (2004).

Balaban et al. (2000) use a simulation approach to assess the consequences of C-5 upgrades. Our approach is methodologically very different from theirs. Balaban et al. do not consider the costs of proposed upgrades, instead focusing on the mission capability consequences.

Nelson et al. (2001) study remanufacturing (i.e., life extension, equipment improvement, and/or mission capability upgrade) to an existing system. Our procedure to analyze modification desirability could apply equally well to an assessment of prospective remanufacturing.

To implement our methodology, we need estimates of the future trajectory of C-5A expenditures ($Expend_{It}$) and availability ($Availability_{It}$). We also need life-cycle estimates of the costs (x) and availability (y) of a replacement system.

We assume, for illustrative purposes, that the C-5A fleet would be replaced by C-17s.

The C-17 is a Boeing-built aircraft that first entered service in the mid-1990s. While smaller than a C-5, the C-17 is more flexible across missions and the airfields it can use. It is also designed to be more reliable and easier to maintain than a C-5, so it can fly more hours per day, on average. See United States Air Force (2004). Figure 3.2 shows a C-17.

Table 3.1 uses information from United States Air Force (2003) and United States Air Force (2004) to compare the aircraft.

We do not know how many C-17s it would take to adequately replace the Air Force's 60 C-5As.¹ Simply as a starting point, we assume one-to-one C-17-for-C-5A replacement. Another possibility is that a different cargo aircraft (the "C-5X") could be developed, but we have no information on such an aircraft's cost or availability characteristics and thus could not use such a new aircraft to illustrate our

¹ Gebman, Batcheler, and Poehlmann (1994) present a detailed analysis of how well C-17s may substitute for the C-5 and other cargo aircraft under different scenarios and assumptions.

Figure 3.2
The C-17 Aircraft



RAND MG241-3.2

Photo from <http://www.af.mil/factsheets/factsheet.asp?fsID=86>.

Table 3.1
Comparing the C-5 and the C-17

| Comparison Point | C-5 | C-17 |
|--------------------------------------|---------|---------|
| First deployed | 1969 | 1993 |
| Crew | 7 | 3 |
| Maximum peacetime takeoff weight, lb | 769,000 | 585,000 |
| Height, feet | 65 | 55 |
| Length, feet | 247 | 174 |
| Wingspan, feet | 223 | 170 |
| Cargo compartment, cubic feet | 36,872 | 19,536 |

methodology. We thus focus on a C-17 alternative to the C-5A, while not endorsing the C-17 relative to other possible replacement approaches.

C-5A and C-17 Costs

To implement our model, we need to estimate both the current levels and future trajectories of C-5A and C-17 costs. We assume a new C-17 has an acquisition cost of \$250 million. United States Air Force (2004) cites a C-17 unit cost of \$236.7 million in FY98 dollars, so we feel \$250 million is a reasonable approximation in FY04 terms.

As was our approach in Keating and Dixon (2003), we separate annual aircraft operating costs for both C-5As and C-17s into personnel and contract maintenance, fuel, reparable, organizational maintenance, aircraft overhaul, engine overhaul, and modifications. As shown in Table 3.2, the Air Force Total Ownership Cost (AFTOC) system provides annual per-aircraft cost factors in these categories for the C-5A, C-5B, and C-17 aircraft.

The surprise in Table 3.2 is that AFTOC's C-5B cost factors are considerably greater than the C-5A cost factors, even though the C-5A is older.

The reason for the AFTOC C-5A/C-5B cost factor difference is that a number of costs common to both systems are allocated in

Table 3.2
AFTOC Annual Per-Active-Aircraft Cost Factors (FY04 \$)

| Cost | C-5A | C-5B | Weighted Average C-5 | C-17 |
|------------------------------------|------------------|-------------------|---------------------------------|-------------------|
| Personnel and contract maintenance | 1,108,646 | 2,060,301 | 1,492,378 | 7,375,834 |
| Fuel | 1,820,978 | 3,110,810 | 2,341,072 | 3,520,102 |
| Reparables | 1,535,123 | 3,760,531 | 2,432,465 | 196,889 |
| Organizational maintenance | 2,855,378 | 7,469,128 | 4,715,761 | 4,826,410 |
| Aircraft overhaul | 647,880 | 1,639,990 | 1,047,924 | 29,979 |
| Engine overhaul | 541,601 | 1,347,720 | 866,649 | 34,153 |
| Mod costs | 821,301 | 700,092 | 772,427 | 1,372,528 |
| Total | 9,330,907 | 20,088,574 | 13,668,675 | 17,355,895 |

AFTOC by flying hours. The typical C-5B flies more hours than the typical C-5A, so a disproportionate share of common costs are allocated to the C-5Bs.² AFTOC per-aircraft cost factors, in other words, do not usefully distinguish between C-5As and C-5Bs for our analytic purposes.

To combat this problem, we used a weighted average C-5 cost factor, as shown in the fourth column of Table 3.2. We formed this weighted average by multiplying the C-5A column by 74 (the number of C-5As in 2002, the year whose data AFTOC used to derive these factors) and the C-5B column by 50 (the FY02 and current number of C-5Bs), then dividing the resultant sum of products by 124.

Table 3.2 also displays AFTOC's C-17 annual per-aircraft cost factors. The C-17 has elevated personnel and contract maintenance costs in conjunction with reduced reparables, aircraft overhaul, and engine overhaul costs compared to the C-5. The C-17's different cost mix emanates from the fact that the C-17 is maintained through Contractor Logistical Support from Boeing, the aircraft's manufacturer. Heretofore, the organic Air Force has provided less maintenance to the aircraft (and Boeing more) than is true for the C-5.

The C-5 has had considerably greater costs per flying hour than the C-17. In FY02, for instance, AFTOC reports total C-5 obligations of about \$1.8 billion and total C-17 obligations of about \$1.3 billion. REMIS, meanwhile, reports FY02 C-5 fleet flying hours of about 88,000 versus about 104,000 C-17 fleet flying hours. Hence, in FY02, the average obligation per C-5 flying hour was about \$20,300 versus \$12,600 for the C-17 (or about \$20,900 versus \$12,900 in FY04 dollars).

The average C-17, however, flew about 1,180 hours in FY02 versus 697 hours per C-5. Hence, the C-17's obligations per aircraft were about \$14.8 million versus \$14.1 million per C-5 (or \$15.2 million versus \$14.5 million in FY04 dollars). These C-17/C-5 obli-

² We appreciate insights from Crash Lively of Battelle on this topic.

gations per aircraft statistics are rough analogs to Table 3.2's AFTOC total cost factors.³

Cost Factor Growth Rates

Given our cost factors, our model next requires estimates of how each factor will grow in the future.

We assume C-5A personnel and contract maintenance will not increase in real terms. However, we assume C-17 personnel and contract maintenance costs will grow at 2 percent per year after the aircraft is 20 years old, because this cost category for the C-17 covers much of the aircraft and engine overhaul expenses (through Contractor Logistical Support) shown to be prone to aging effects. See Pyles (2003).


We reprise assumptions from Keating and Dixon (2003) that fuel costs grow 0.6 percent per annum as aircraft age and reparable costs grow at 3.5 percent per annum. The 0.6 percent fuel cost growth rate assumption originally came from Stoll and Davis (1993).

Organizational maintenance is maintenance done at the installation level. Figure 3.3 shows C-5A installation-level on- and off-equipment maintenance hours per possessed (i.e., not held by the depot system) hour since the beginning of FY00 according to REMIS. (If an aircraft is not possessed, i.e., it is held by the depot system, it should not receive installation-level on- and off-equipment maintenance.)

³ Our unit of analysis is the aircraft, not the flying hour, thus we use the annual per-aircraft cost factors from Table 3.2 in our analyses. A typical C-17 flies more hours than a typical C-5, but, as shown in Table 3.1, a C-17 is a smaller aircraft that carries smaller loads. In this work, we are agnostic as to the number of C-17s that would be needed to replace the C-5A fleet, beyond starting with a one-for-one replacement rate baseline. In Figure 3.8, later in this chapter, we show how the optimal timing of C-5A replacement varies as a function of the number of C-17s needed to replace the 60 C-5As.

Figure 3.3

C-5A On- and Off-Equipment Maintenance Hours Per Possessed Hour




RAND MG241-3.3

There is no evidence in Figure 3.3 of any adverse C-5A aging effect in terms of on- and off-equipment maintenance. Indeed, installation-level maintenance hours per possessed hour have fallen. In our model application, we assume C-5A organizational maintenance costs will be static in real terms into the future. We make the same assumption for the C-17.

Aircraft overhaul, meanwhile, occurs at the depot level. The C-5A has had considerable problems with increasing depot-level maintenance hours per aircraft, as shown in Figure 3.4. These Air Force data are described in Pyles (2003). The C-5A's per-aircraft programmed depot maintenance (PDM) package has grown at a 9.6 percent annual rate, in maintenance hours, over the period of these data. Therefore, we assume C-5A aircraft overhaul costs will grow at 9.6 percent per year into the future.

Figure 3.4
C-5A PDM Hours Per Aircraft



RAND MG241-3.4

There is considerable uncertainty about the future trajectory of PDM hours and, consequently, costs for the C-5A and other aircraft. Some experts feel that a plateau of PDM hours will soon be achieved. At the same time, we have not seen conclusive evidence of such plateaus. It is also unclear whether the observed rate of growth in maintenance hours per aircraft is a reasonable proxy for growth in aircraft overhaul costs per aircraft (which include nonlabor factors such as materiel used in PDM). As we suggested in Keating and Dixon (2003), it is important to undertake robustness analysis to assess, for instance, how and whether results change based on one's beliefs as to the future evolution of PDM costs. (We display such a robustness exploration for the aircraft overhaul cost growth rate later in Figure 3.10.)

We reprise Keating and Dixon (2003) and assume 3.5 percent growth in engine overhaul costs. We assume the RERP will cut both C-5A annual aircraft and engine overhaul costs by 10 percent.

We also reprise Keating and Dixon (2003) in assuming 2.1 percent growth in modification costs, separate from the \$75 million C-5A RERP. We further assume a C-17 will get a comparable RERP in its 44th year of operation (if it is operated that long).

Table 3.3 reprises Table 3.2's AFTOC weighted average C-5 and C-17 parameters and shows our assumptions about their future growth rates. In general, we assume that the C-17 will manifest aging effects similar to those seen on the C-5A but that many of these effects will not start for 15–20 years.

Table 3.3
Annual Sustainment Costs and Assumed Growth Rates of the C-5A and C-17 Fleets (per aircraft)


| Category | C-5A Millions FY04\$/Tail | C-5A Assumed Growth Rate | C-17 Millions FY04\$/Tail | C-17 Assumed Growth Rate |
|------------------------------------|-----------------------------------|-----------------------------------|--|---------------------------------------|
| Personnel and contract maintenance | 1.49 | 0 | 7.38 | 0 for first 20 years, 2% thereafter |
| Fuel | 2.34 | 0.6% | 3.52 | 0.6% |
| Reparables | 2.43 | 3.5% | 0.20 | 3.5% |
| Organizational maintenance | 4.72 | 0 | 4.83 | 0 |
| Aircraft overhaul | 1.05 | 9.6%; RERP cuts cost 10% | 0.03 | 0 for first 15 years, 9.6% thereafter |
| Engine overhaul | 0.87 | 3.5%; RERP cuts cost 10% | 0.03 | 0 for first 20 years, 3.5% thereafter |
| Modifications | 0.77 plus \$75 RERP in 2015 | 2.1% | 1.37 plus \$75 RERP when 44 years old (if optimal) | 0 for first 15 years, 2.1% thereafter |
| Total | 13.67 | | 17.36 | |

Aircraft Availability

Our model also needs estimates of the incumbent's and replacement's aircraft availability levels ($Availability_{It}$ and $Availability_{Rt}$).

Figure 3.5 shows the C-5A's monthly mission-capable (MC) rate since January 1991 as well as what we term the aircraft's composite availability rate. These data are from REMIS. We define the composite rate as the product of the aircraft's MC rate and its possessed rate. We feel the composite availability rate is a more useful metric than the MC rate because the composite rate reflects the rate at which field-level commanders can actually use the aircraft. The composite availability rate tallies the fraction of the total fleet that is both possessed by the operating commands and mission capable. The MC rate, by contrast, is defined only for the fraction of aircraft that is

Figure 3.5
C-5A MC and Composite Availability Rates



possessed by operating commands; aircraft in the depot system are not considered.

Both the C-5A composite availability and MC rates have fallen somewhat as the aircraft has aged, but the magnitude of the decline has not been enormous. Over the period of these data, the composite availability rate has trended downward at about 1.1 percent per year; the mission-capable rate has trended downward at about 0.8 percent per year. As shown in Figure 3.5, there has been considerable month-to-month variation in both rates. The last few years have had more favorable composite availability rates than the years preceding them.

We assume, as a starting point, that the \$75 million mod will increase C-5A composite availability 50 percent, from the low 40 percent range to around 60 percent.


We use actual 1999–2002 data on C-17 composite availability from REMIS to estimate the first four years of a new C-17's composite availability, and then project it to decline at the C-5A's 1.1 percent rate.

Optimal C-5A Retirement Calculation

Bringing together our various parameters and assumptions, Figure 3.6 shows our baseline C-5A analysis. The broken line is the estimated incremental cost per available year of the C-5A; the horizontal line is the C-17's estimated average cost per available year (about \$46 million). The large spike at age 44 represents the \$75 million C-5A mod that will remove an aircraft undergoing the mod from the fleet for six months. The C-5A incremental cost line is lower after the mod than before, reflecting the projected post-mod increase in composite availability and reduction in costs.

In our baseline case, our model indicates that it is not worthwhile to modify the C-5A in 2015 at age 44. The annualized cost per available year is lower if the Air Force retires the aircraft at age 43 in front of the mod ($AC_{NM} = \$42.7$ million) rather than if the aircraft is retired at age 53, nine years after the RERP ($AC_M = \$44.3$ million). See Table 3.4.

Figure 3.6
C-5A Optimal Replacement Baseline Case



RAND MG241-3.6

Robustness Explorations

Figure 3.6 and Table 3.4 are by no means the end of the story. Our model implementation includes a number of important assumptions that have the potential to affect the analytic results. We therefore assess which parameters were particularly influential in indicating whether it is worthwhile to undertake the mod program.

One key parameter is the timing of the RERP. In our baseline case, we assume the median C-5A gets the RERP in 2015. Holding our other parameters fixed, Figure 3.7 shows that the RERP is


Table 3.4
Estimated C-5A Costs

| C-5A Median Age | Year | $Expend_{It}$ (Millions of FY04\$) | $Availability_{It}$ (%) | Incremental Cost (Millions of FY04\$) | Annualized Cost (Millions of FY04\$) |
|-----------------------|------|--|----------------------------|---|--|
| 40 | 2011 | 16.5 | 42.2 | 39.1 | 43.0 |
| 41 | 2012 | 16.9 | 41.7 | 40.6 | 42.9 |
| 42 | 2013 | 17.4 | 41.2 | 42.1 | 42.8 |
| 43 | 2014 | 17.9 | 40.8 | 43.8 | 42.7 |
| 44 | 2015 | 93.4 | 20.2 | 462.9 | 45.9 |
| 45 | 2016 | 18.4 | 59.8 | 30.7 | 45.6 |
| 46 | 2017 | 18.9 | 59.2 | 32.0 | 45.3 |
| 47 | 2018 | 19.5 | 58.5 | 33.4 | 45.0 |
| 48 | 2019 | 20.2 | 57.9 | 34.8 | 44.8 |
| 49 | 2020 | 20.8 | 57.3 | 36.4 | 44.7 |
| 50 | 2021 | 21.6 | 56.6 | 38.1 | 44.5 |
| 51 | 2022 | 22.3 | 56.0 | 39.9 | 44.4 |
| 52 | 2023 | 23.2 | 55.4 | 41.8 | 44.3 |
| 53 | 2024 | 24.1 | 54.8 | 43.9 | 44.3 |
| 54 | 2025 | 25.0 | 54.2 | 46.2 | 44.3 |
| 55 | 2026 | 26.1 | 53.6 | 48.6 | 44.4 |

worthwhile (and the optimal retirement age is 54, corresponding to 2025) if the median C-5A receives the RERP in 2010 or earlier. (In Figures 3.7–3.10, we have placed the vertical axis at the parameter level we assumed in the base case.) In our baseline case, the RERP is scheduled for 2015 when the median C-5A is 44 years old, so any computed optimal retirement year greater than 44 suggests that the RERP should be undertaken.

Figure 3.7

C-5A Optimal Replacement as a Function of When the Median C-5A Gets the RERP




RAND MG241-3.7

Intriguingly, Figure 3.7's result echoes an observation by General John W. Handy, commander of U.S. Transportation Command and Air Mobility Command. Tirpak (2004) notes "By 2012," Handy said, "it may be that the whole notion [of performing the RERP on the C-5As] is overcome by events. The A models might be too far gone to be worth the investment." An option we have not explored but suggested by this calculation would be to perform the RERP on the C-5A fleet before it is done on the C-5B fleet.

Another key parameter is the number of C-17s it would require to replace the existing fleet of 60 C-5As. In Figure 3.6, we assume C-17s could replace C-5As one-for-one. However, as shown in Figure 3.8, if it takes 70 or more C-17s to acceptably replace the 60 C-5As,

Figure 3.8
C-5A Optimal Replacement as a Function of the Number of C-17s Required




our model indicates that the RERP should be undertaken in 2015 and the C-5A fleet should not be retired until a later date.

In the base case, we assumed the C-5A RERP increases the C-5A's composite availability by 50 percent, i.e., from the low 40 percent range to around 60 percent. Tirpak (2004) notes that modified C-5As will have a mission-capable goal of 75 percent; Lockheed Martin goes so far as to project a modified C-5A departure reliability rate of nearly 95 percent.

If the C-5As' operating command possession rate were 85 percent, a 75 percent mission-capable rate would result in a 64 percent composite availability rate, while a 95 percent mission-capable rate would result in an 81 percent composite availability rate. An 81 percent composite availability rate would nearly double the composite availability rates we project for the early part of the 2010 decade (in

Figure 3.9

C-5A Optimal Replacement as a Function of the Composite Availability Gain After the Age 44 Mod




RAND MG241-3.9

the low 40 percent range). As shown in Figure 3.9, our model indicates that a composite availability gain on that scale would make a 2015 RERP worthwhile, holding our other parameters fixed.

Another parameter of interest is the growth rate in aircraft overhaul costs, a topic initially presented in Figure 3.4. In Figure 3.10, we show that our model indicates the 2015 RERP is worthwhile if the aircraft overhaul cost growth rate decreases to 4 percent or less, rather than to our baseline 9.6 percent assumption.

Figure 3.10
C-5A Optimal Replacement as a Function of the Aircraft Overhaul Cost Growth Rate Past Age 31



RAND MG241-3.10

Discussion

In this chapter, we first illustrated our methodology for how the desirability of a mod program can be assessed. We then used our methodology to identify several key issues in C-5A decisionmaking.

One key issue is the timing of the proposed C-5A RERP. In accord with statements by Air Force leadership, our model indicates that the C-5A RERP is more desirable if it comes sooner, other things equal.

On the other hand, an important unknown is how much of an increase in composite availability will emanate from the RERP. Tirkpak (2004) notes that the Air Force is testing out the RERP on two C-5Bs and a C-5A to determine if the improvements will provide the anticipated boost in performance and reliability. Thus, it appears that the Air Force faces a tradeoff between the value of better information about the RERP and the cost of delaying implementation of the program.

Another issue is what system would replace the C-5A and in what proportion. If the C-17 is the C-5A's replacement, our model suggests that the RERP is not optimal, given other parameters, if fewer than 70 C-17s can replace the 60 C-5As. As noted by Gebman, Batchelder, and Poehlman (1994), there are widely varying views as to the appropriate C-17/C-5A replacement ratio. Another possibility is that a different "C-5X" aircraft would be better—perhaps more capable and/or cheaper—than C-17s. A better replacement aircraft would tend to encourage replacement of, and discourage doing the RERP on, the C-5A.

Our results should be viewed as illustrative and suggestive, rather than definitive, in light of our many assumptions and the shortcomings of the C-5A cost factors.

Consideration of C-5A Depot-Level Capacity

In our original analysis, highlighted in Chapter Two, we presented a methodology to assess the optimal time at which to retire an aging system. Chapter Three illustrated how, in an extension of our methodology, the choice between retiring and modifying an aging system might be assessed.

In this chapter, we illustrate how our methodology might be further extended to broader resource-allocation decisions. For example, an alternative to retiring an aging system is to invest more resources in its maintenance. For the sake of parsimony, we will speak of investments in “depot-level capacity.” Depot-level capacity, however, is meant as a general rubric for any steps that could be taken to get aircraft out of maintenance more quickly and available to operators more often. Examples of investment in depot-level capacity may include (but are not limited to):

- Building additional maintenance facilities at Air Force-owned and/or private sector depots.
- Purchasing additional repair equipment for use at installations and/or depots.
- Hiring more and/or better labor to maintain aircraft.
- Paying overtime to workers to get aircraft out of maintenance faster.
- Buying additional spare parts to reduce delays from waiting for parts.

We will not attempt here to assess the optimal form of investment in additional depot-level capacity (for example, on the margin, is it better to have more spare parts or more repair equipment?). We will also avoid issues related to the tradeoffs between private-sector maintenance and maintenance at government-owned facilities.

Our aspirations in this chapter are more modest. We start with a thought experiment: In Chapter Three, we estimated that the appropriately annualized cost of an available year of a C-17 would be about \$46 million. We further assumed that the Air Force would someday be willing to pay this price. Hence, by assumption, an available C-17 year must be worth at least \$46 million to the Air Force.

We further assumed, in our baseline case, that the Air Force trades off C-5As and C-17s one-to-one. We also assume here that the Air Force's preferences are linear, that is, not subject to diminishing marginal utility. A 10 percentage point increase in composite availability is worth the same whether starting from a base of 40 percent or 80 percent.

Stringing together our assumptions, we infer that an available C-5A year is also worth at least \$46 million to the Air Force.

This inference, which we readily admit is subject to considerable controversy, is interesting as it pertains to depot-level capacity choice. Specifically, given a belief (or inferred preference) about how much an available aircraft year is worth to the Air Force, we can calibrate the desirability of increases in depot-level capacity that put aircraft in operators' hands faster. If depot-level capacity increased, aircraft could spend less time in depot-level maintenance and repair (e.g., shorter or fewer queues in the depot process), so composite availability could be increased.

We spoke with personnel who provide depot-level maintenance of aircraft. We learned that sizable chunks of time in the depot process are spent in queues. For example, an aircraft might have to wait for a test stand to free up from an aircraft ahead of it. Or an aircraft might be found to have a broken part that must be ordered through the supply system. Perhaps key laborers are working on a different aircraft, so delay ensues. We therefore developed a depot-level queuing model in which changes in depot-level capacity can be used to

change the magnitude of queuing. With more capacity, queues are shorter (or vanish) and aircraft move through the process faster and return to operating commands sooner.

Because depot-level capacity is costly, it would presumably be prohibitively expensive to have no depot-level queuing. However, given the inferred value of aircraft, are current levels of depot-level queuing optimal?

We model the PDM process as a multinode or multistep closed network. Building upon Zahorjan et al. (1982), let us define:

- R_k = mean residence time in days at node k , including repair and queuing time; R_k is a function of N , the total number of aircraft.
- L_k = hands-on repair time in days at node k .
- n_k = mean queue length at node k .
- c_k = number of servers at node k . There can be more than one server at a given node. The number of servers is our measure of depot-level capacity.
- K = total number of nodes in the depot process. There are K steps in putting a given aircraft through PDM.

$$\bullet R_0 = \sum_{k=1}^K L_k$$

= minimum hands-on flow time in days to complete PDM,
with no queuing.

A key relation used in the Zahorjan et al. (1982) model is the mean value analysis equation, $R_k = (1 + n_k)L_k$, to estimate the residence time R_k at node k . The expression indicates that the residence time is the sum of the service and queueing times of the N th customer, the latter given by the average queue length (with $N - 1$ customers in the network) times the service time. The Zahorjan et al. analysis assumes that there is only one server at node k .

For the more general case with c_k servers at node k , the time the N th customer waits in queue is more complicated. There are three cases: (1) if there are $n < c_k$ customers at node k , the arriving customer has no wait; (2) if there are $n = c_k$ customers at node k , the customer must wait in queue for one customer to be served at rate c_k / L_k ; and (3) if there are $n > c_k$ customers at node k when the N th customer arrives, the newly arrived customer must wait in queue until $n - c_k + 1$ customers are serviced at rate c_k / L_k . This gives

$$R_k = [1 + \overline{n_k} + \sum_{n=0}^{c_k-2} (c_k - 1 - n) p_k(n | N - 1)] \frac{L_k}{c_k},$$

where $p_k(n | N - 1)$ is the marginal probability that n customers are at node k when there are $N - 1$ customers in the network. See Gross and Harris (1998). This probability can be calculated recursively, but that defeats the simplicity of the Zahorjan et al. (1982) model.

There exist bounds on R_k . Using the constraint that $p_k \leq 1, \forall n$ and that

$$\sum_{n=0}^{c_k-2} p_k(n | N - 1) \leq 1,$$

it follows that

$$\sum_{n=0}^{c_k-2} (c_k - 1 - n) p_k(n | N - 1) \leq c_k - 1$$

which gives the bound

$$R_k < [1 + \overline{n_k} + (c_k - 1)] \frac{L_k}{c_k},$$

which can be reexpressed as

$$R_k < L_k + \frac{\overline{n_k L_k}}{c_k}.$$

L_k is the actual hands-on repair time; $(n_k L_k) / c_k$ is mean time spent in a queue before repair commences. Hence,

$$\sum_{k=1}^K R_k < R_0 + \sum_{k=1}^K \frac{\overline{n_k L_k}}{c_k}.$$

The actual time in the depot system is composed of hands-on repair time plus time spent in queue.

There are a number of assumptions lying behind Zahorjan et al.'s and our use of this inequality:

- Zahorjan et al. used Little's Formulas for queuing. Little's Formulas do not rely on distributional assumptions about workload arrival or service times. However, Little (1961) requires strict stationarity¹ and a metrically transitive (ergodic) arrival process. Gross and Harris (1998) present an overview of Little's Formulas.
- Zahorjan et al.'s approach assumes the network is separable, i.e., work from multiple steps cannot occur simultaneously.² Also, one cannot divert an aircraft queued at one node to another node that happens not to have a queue. Each node in the network behaves as if it was a single queuing system. See Bolch et al. (1998) for a discussion of separability.
- We assume the network is balanced, i.e., the relative utilization of all nodes is the same. Bolch et al. (1998) also discuss network

¹ Box, Jenkins, and Reinsel (1994), p. 24, note "for a discrete process to be strictly stationary, the joint distribution of any set of observations must be unaffected by shifting all the times of observation forward or backward by any integer amount k ."

² Dietz and Jenkins' (1997) analysis allows for multiple types of repair to be performed concurrently by viewing the aircraft as generating temporary clones that are rejoined into a single entity when all repairs are complete. We do not use their approach in our analysis.

balance. Any addition of capacity will be accomplished in such a way as to maintain a balance in the workloads. A balanced network is the most efficient type of network in terms of throughput and resource utilization.

The Air Force is interested in the fraction of its fleet that is possessed by operating commands—not tied up in depot-level repair. Suppose an aircraft is on a five-year PDM cycle: The aircraft spends

$$\sum_{k=1}^K R_k$$

days in PDM then five years operating.³ On average, then, the command possession rate would be

$$CP = \frac{5 * 365.25}{5 * 365.25 + \sum_{k=1}^K R_k}.$$

Re-arranging terms, we find

$$\sum_{k=1}^K R_k = 1826.25 * \left(\frac{1}{CP} - 1 \right).$$

Then if we plug into the inequality above, we find

$$CP > \frac{1826.25}{1826.25 + R_0 + \sum_{k=1}^K \frac{n_k L_k}{c_k}}.$$

³ With a five-year PDM cycle, an aircraft always spends five years out of the depot system, no matter how long the preceding stay in the depot system took. Five years is measured from time of discharge from the depot system to time of next induction into the depot system.

The number 1826.25 represents five years of operating command possession days between PDM visits, R_0 is hands-on maintenance time during PDM and repair, and

$$\sum_{k=1}^K \frac{\overline{n_k L_k}}{c_k}$$

is time spent in queues during PDM.

As an example, let us again consider the C-5A, an aircraft on a five-year PDM cycle. Suppose the C-5A's $R_0 = 108$ days. (We will view R_0 as exogenous in this exploration. It also includes time spent in non-PDM depot-level repair.) In FY02, the observed C-5A command possession rate was about 0.829. Our inequality is then

$$\frac{1826.25}{1826.25 + 108 + \sum_{k=1}^K \frac{\overline{n_k L_k}}{c_k}} < 0.829,$$

which implies

$$\sum_{k=1}^K \frac{\overline{n_k L_k}}{c_k} > 268.$$

Next, consider the effect of increasing capacity, e.g., purchasing more repair equipment, spare parts, or hiring more workers. Obviously, the effect on queuing time,

$$\sum_{k=1}^K \frac{\overline{n_k L_k}}{c_k},$$

and hence on the command possession rate, CP , will depend on which specific node k 's receive the extra capacity.

As a simple illustrative case, let $K = 2$. Let the numerator of queuing time at node one $n_1 L_1 = a_1$ and the numerator of queuing time at node two $n_2 L_2 = a_2$. Without loss of generality, assume $a_2 > a_1$. Suppose total capacity available $c_1 + c_2 = Y$. How should the total server capacity be optimally distributed across the two nodes so as to minimize the sum of mean queuing time?

Let X denote the amount of capacity devoted to node 1 and $Y - X$ the amount of capacity on node 2. Then the sum of mean queuing time we wish to minimize would be

$$\frac{a_1}{X} + \frac{a_2}{Y - X}.$$

We differentiate this term with respect to X and then set the derivative equal to zero, resulting in the equation

$$\frac{a_1}{X^2} = \frac{a_2}{(Y - X)^2}.$$

Solving for X , we find the relevant root

$$X = Y * \left(\frac{\sqrt{a_1 a_2} - a_1}{a_2 - a_1} \right) \text{ and } Y - X = Y * \left(\frac{a_2 - \sqrt{a_1 a_2}}{a_2 - a_1} \right).$$

With these values of X and $Y - X$, the sum of mean queuing time would be

$$\frac{1}{Y} \left(\frac{(a_2 - a_1)^2}{a_1 + a_2 - 2\sqrt{a_1 a_2}} \right).$$

(In the degenerate case of $a = a_1 = a_2$, $X = Y - X = Y / 2$ and the sum of mean queuing time would be $4a / Y$.)

The important point in this exercise is to observe that the sum of mean queuing time is a function of the two nodes' numerators of

queuing times, a_1 and a_2 , that we will label M ($M = f(a_1, a_2)$), divided by Y , the total capacity at the two nodes:

$$\sum_{k=1}^K \frac{\overline{n_k L_k}}{c_k} = \frac{M}{Y}.$$

To illustrate this phenomenon, suppose

$$\sum_{k=1}^K \frac{\overline{n_k L_k}}{c_k} = 268.$$

Suppose, too, that current total C-5A capacity is 12. Then M , the numerator in the sum of mean queuing time formula M/Y , equals 3214. Then our command possession inequality would be

$$CP > \frac{1826.25}{1826.25 + 108 + \frac{3214}{Y}}.$$


If capacity then doubled to $Y = 24$, our new command possession lower bound would be 0.883, up from 0.829 with half as much depot-level capacity. Figure 4.1 plots the lower bound on the command possession rate as a function of C-5A depot-level capacity.

⁴ This formula also holds for cases of $K > 2$. For instance, if $K = 3$ and $\overline{n_1 L_1} = a_1$, $\overline{n_2 L_2} = a_2$, and $\overline{n_3 L_3} = a_3$, then

$$\sum_{k=1}^3 \frac{\overline{n_k L_k}}{c_k} = \frac{(1 + \sqrt{\frac{a_2}{a_1}} + \sqrt{\frac{a_3}{a_1}})}{Y} * (a_1 + \frac{a_2}{\sqrt{\frac{a_2}{a_1}}} + \frac{a_3}{\sqrt{\frac{a_3}{a_1}}}) = \frac{M}{Y},$$

where $M = f(a_1, a_2, a_3)$.

Figure 4.1
Lower Bound on C-5A Operating Command Possession Rate as a Function of Depot-Level Capacity



Asymptotically, with unlimited capacity, there is no queuing and the command possession rate lower bound is

$$\frac{1826.25}{1826.25 + R_0}.$$

With $R_0 = 108$, the infinite depot-level capacity operating command possession rate would be about 0.944.

Growing Hands-On Time During PDM and Repair

In Figure 3.4, we presented evidence of considerable historical increases in C-5A PDM hours per aircraft. Based on that figure, we assumed aircraft overhaul costs would continue to rise at 9.6 percent

per year into the future. It is reasonable to think the C-5A's R_0 , the minimum possible flow time, will also grow as these aircraft age, though perhaps not at this specific rate.

If R_0 grows, the operating command possession rate will fall, holding capacity constant. In Figure 4.2, we plot the lower bound on the C-5A command possession rate for different levels of R_0 , holding capacity at 12.

As the minimum hands-on maintenance and repair time R_0 grows, our model implies that the level of capacity required to maintain a specific rate of operating command possession increases. In Figure 4.3, we show the required level of C-5A depot-level capacity to maintain an 80 percent command possession rate for different levels of R_0 . The required level of capacity grows nonlinearly in R_0 .

Figure 4.2
Lower Bound on C-5A Operating Command Possession Rate as a Function of R_0





Figure 4.3
Required Level of C-5A Capacity as a Function of Minimum Hands-On Time, R_0



RAND MG241-4.3

Figure 4.3's nonlinearity illustrates that it would be very worrisome and expensive if R_0 grows considerably.

Consideration of R_0 and depot-level capacity changes our repair-replace analysis framework. Specifically, one no longer uses a historical trend for the composite availability rate. Instead, while one continues to use an MC rate trend, the possession rate is a function of R_0 and capacity.

For illustrative purposes, we assume R_0 grows at 2 percent per year. We also assume the total number of C-5As drops to 60 in FY07, as currently planned by the Air Force. (The effect of the decline in the number of C-5As is to increase the amount of depot-level capacity that exists, relative to the fleet size.) With these assumptions, our model indicates an optimal retirement age of 43, just in front of the RERP, the same as indicated in Figure 3.6. (While the answers coincided in this case, one will not always have the same estimated retirement age when depot-level capacity is considered. Our assumptions about R_0 growth, for instance, can lead to different projections of future composite availability than the age-driven extrapolation we used in Chapter Three.)

Investing in More Depot-Level Capacity

Extending this illustration, we explore the desirability of investment in additional C-5A depot-level capacity. We arbitrarily assume each unit of extra capacity would have a one-time cost of \$25 million and an annual incremental cost of \$2 million for every year until the C-5A is retired.⁵

We assume a one-time choice of whether or not to purchase additional C-5A depot-level capacity in 2005.⁶ Again, we assume R_0 grows at 2 percent per year. The decision, then, is whether to invest in more capacity and, if so, how much more.

We also assume we could use FY02 parameters (e.g., the 82.9 percent operating command possession rate) to characterize the C-5A


⁵ This one-time plus annual incremental cost function is similar to the approach of Cakanyildirim, Roundy, and Wood (2004) in their analysis of a production facility deciding whether to buy extra machines.

⁶ It would be a considerably more complex exercise, beyond the illustrative aspirations of this chapter, to make depot-capacity choice a dynamic decision one could revisit annually. We have also not considered the possibility of increased depot capacity for the C-17 or other replacement aircraft. The option of investment in increased depot capacity for the replacement aircraft could make it more desirable to retire the incumbent aircraft. Here we assume the increased depot-level capacity affects only the C-5A.

depot system (beyond R_0 growth and potentially changing capacity) in future years.

With the parameters we assume, our model finds that it would be desirable to undertake investment in additional depot-level capacity, as shown in Figure 4.4. Figure 4.4’s vertical axis is the annualized cost per available year (AC_{NM}) of the C-5A and its eventual replacement C-17s, including whatever costs are borne to invest in additional C-5A depot-level capacity.


Figure 4.4
 Annualized Cost Per Available Year with Investment in Extra C-5A Depot-Level Capacity



With the assumptions described above, it is optimal to invest in five extra units of C-5A depot-level capacity (indeed, buying up to 12 extra units is found to be better than none). Five extra units represent more than a 40 percent increase in C-5A depot-level capacity. Yet, the model indicates that the optimal C-5A retirement age does not change. It remains optimal to retire the C-5A in front of the RERP in 2015, holding fixed our other parameters.

While the optimal retirement age does not, with these parameters, change, investing in the optimal five extra units of C-5A depot-level capacity increases the aircraft's operating command possession rate. See Figure 4.5. In both cases, the C-5A's estimated possession

Figure 4.5
C-5A Possession Rate With and Without Extra Depot-Level Capacity



rate increases from ages 34 to 36 (2005–2007), reflecting the retirement of 12 of the aircraft. From age 36 (2007) forward, the new steady-state of 60 C-5As is achieved and the composite availability rate drifts downward, reflecting assumed growth in R_0 and further decline in the mission-capable rate.

It is interesting that, with our parameters, extra depot-level capacity, while desirable, does not increase the optimal C-5A retirement age.

We think our findings are explained by the inferred preference exercise upon which this analysis is built: We assume the Air Force will eventually replace the C-5A with the C-17, a nontrivially expensive aircraft. Hence, we infer that the Air Force puts a high valuation on an available C-17, and hence C-5A, year. Given this inferred preference, the model is then averse to having large-scale queuing and delay in the C-5A depot maintenance system. Extra C-5A depot-level capacity is recommended, even when it does not lengthen the aircraft's life. Given the inferred value of these aircraft, it is not reasonable to have them wait in many or long depot-level queues.

As noted in Chapter Three, another possibility is a “C-5X” that could be cheaper and/or better than the C-17. If the alternative to the C-5A is more desirable, the inferred value of additional C-5A depot-level capacity will drop. This chapter's calculation relies heavily on the value inference implicit in the belief that the Air Force is willing to pay for C-17s to replace C-5As.

Of course, the results in this chapter are highly speculative. The other side of the calculation, which we have simply parameterized here, is how much it would cost to increase depot-level capacity and reduce queuing in the system. Another study could assess the most cost-effective ways to increase depot-level capacity. Our exploration, however, is suggestive that such an increase could be desirable, based on the inference that available cargo aircraft years are quite valuable to the Air Force.

Conclusions

This study has expanded on the repair-replace methodology we set forth in Keating and Dixon (2003). In that report and in Chapter Two, we presented a methodology for determining when it is optimal to replace an aging system with a new aircraft.

In Chapter Three of this paper, we studied a related problem: Should an aging system receive a substantial modification or should it be retired prior to receiving this modification? We used data on the C-5A cargo aircraft to illustrate implementation of our model. Based on our data and assumptions, we found that the large-scale modification proposed for the C-5A is desirable if it happens soon, but it becomes less desirable as it moves further into the future. Other key parameters include the magnitude of C-5A composite availability improvement after the mod as well as the ratio at which C-17s (the assumed replacement aircraft) can replace C-5As. We caution, however, that our C-5A cost parameters are problematic, so results should be viewed only as illustrative.

In Chapter Four, we presented a new methodology to assess the desirability of additional investment in depot-level capacity (where such capacity investment could take the form of additional facilities, equipment, labor, and/or spare parts).

Our Chapter Four exercise was built on a string of inferred preferences: Because the Air Force is eventually willing to pay the cost of a replacement system, the cost per available year of that new system can be argued to be a lower bound on how much the Air Force values

available aircraft. Given this inferred valuation, it might then be reasonable to invest in depot-level capacity for an existing system so as to reduce the time that the system spends delayed in depot-level maintenance and repair queues.

Applying this methodology to the C-5A (with many assumptions about key parameters), we found potential benefit in increased C-5A depot-level capacity without, however, extending the aircraft's optimal life. We caution that this result is highly speculative. It relies on numerous assumptions such as linearity in the Air Force's utility function for available aircraft. It may also be driven by the illustrative analysis of the C-17 as the C-5A's prospective replacement.

This exploratory study of methodology did not investigate how best to increase depot-level capacity. Instead, it simply took preliminary steps toward assessing the potential desirability of depot-level capacity increases. It appears to be suboptimal to have large-scale queuing, and hence delay, in the depot system for an aircraft whose availability the Air Force values highly.

References

- Balaban, Harold S., Robert T. Brigantic, Samuel A. Wright, and Anthony F. Papatyi, "A Simulation Approach to Estimating Aircraft Mission Capable Rates for the United States Air Force," in J. A. Joines, R. R. Barton, K. Kang, and P. A. Fishwick (eds.), *Proceedings of the 2000 Winter Simulation Conference*, <http://www.informs-cs.org/wsc00papers/136.PDF>, 2000.
- Bolch, Gunter, Stefan Greiner, Hermann de Meer, and Kishor S. Trivedi, *Queuing Networks and Markov Chains*, New York: John Wiley & Sons, Inc., 1998.
- Box, George E.P., Gwilym M. Jenkins, and Gregory C. Reinsel, *Time Series Analysis: Forecasting and Control*, Englewood Cliffs, NJ: Prentice Hall, Third Edition, 1994.
- Cakanyildirim, Metin, Robin O. Roundy, and Samuel C. Wood, "Optimal Machine Capacity Expansions with Nested Limitations under Stochastic Demand," *Naval Research Logistics*, Vol. 51, No. 2, March 2004, pp. 217–241.
- Dietz, Dennis C., and Richard C. Jenkins, "Analysis of Aircraft Sortie Generation with the Use of a Fork-Join Queuing Network Model," *Naval Research Logistics*, Vol. 44, No. 2, March 1997, pp. 153–164.
- Gebman, Jean R., Lois J. Batchelder, and Katherine M. Poehlmann, *Finding the Right Mix of Military and Civil Airlift: Issues and Implications*, Santa Monica, CA: RAND Corporation, MR-406-AF, 1994.

- Greenfield, Victoria A., and David M. Persselin, *An Economic Framework for Evaluating Military Aircraft Replacement*, Santa Monica, CA: RAND Corporation, MR-1489-AF, 2002.
- Gross, Donald, and Carl M. Harris, *Fundamentals of Queueing Theory*, New York: John Wiley & Sons Inc., Third Edition, 1998.
- Jablonski, David A., "Air Force Fleet Viability Board Releases C-5A Assessment," *Air Force Print News Today*, <http://www.af.mil/news/story.asp?storyID=123008163>, July 15, 2004.
- Keating, Edward G., and Matthew Dixon, *Investigating Optimal Replacement of Aging Air Force Systems*, Santa Monica, CA: RAND Corporation, MR-1763-AF, 2003.
- Little, John D.C., "A Proof for the Queuing Formula: $L = \lambda W$," *Operations Research*, Vol. 9, No. 3, May–June 1961, pp. 383–387.
- Nelson, Richard, et al., "Remanufacture of Defense Systems: Case Studies of Cost Experiences and Lessons for Future Programs," Department of Defense Working Paper, November 2001.
- Office of Management and Budget, *Circular No. A-94: Guidelines and Discount Rates for Benefit-Cost Analysis of Federal Programs*, <http://www.whitehouse.gov/omb/circulars/a094/a094.html>, October 1992 and February 2004 (Appendix C).
- Pike, John, "C-5 Galaxy," *GlobalSecurity.org*, <http://www.globalsecurity.org/military/systems/aircraft/c-5.htm>, November 2001.
- Pyles, Raymond A., *Aging Aircraft: USAF Workload and Material Consumption Life Cycle Patterns*, Santa Monica, CA: RAND Corporation, MR-1641-AF, 2003.
- Stoll, Laurence, and Stan Davis, *Aircraft Age Impact on Individual Operating and Support Cost Elements*, Patuxent River, MD: Naval Maintenance Office Logistics Engineering Department, Resource Analysis Division, July 1993.
- Tirpak, John A., "Saving the Galaxy," *Air Force Magazine*, Vol. 87, No. 1, <http://www.afa.org/magazine/Jan2004/0104galaxy.asp>, January 2004, pp. 30–35.
- United States Air Force, *C-5 Galaxy* fact sheet, <http://www.af.mil/factsheets/factsheet.asp?fsID=84>, August 2003.

- United States Air Force, *C-17 Globemaster III* fact sheet, <http://www.af.mil/factsheets/factsheet.asp?fsID=86>, June 2004.
- Zahorjan, John, Kenneth C. Sevcik, Derek L. Eager, and Bruce Galler, “Balanced Job Bound Analysis of Queuing Networks,” *Communications of the ACM*, Vol. 25, No. 2, February 1982, pp. 134–141.